

Chapter 5: Overview of the Enstore Servers

In this chapter we describe the software modules that act as Enstore servers and the libraries with which they interact. The servers include:

- File Clerk (FC)
- Volume Clerk (VC)
- Library (also called a Robot) and Virtual Library (VL)
- Library Manager (LM)
- Mover (MV)
- Media Changer (MC)
- Configuration Server (CS)
- Inquisitor
- Alarm Server (AS)
- Log Server (LS)
- Event Relay (ER)

5.1 File Clerk

The File Clerk (FC) is a server that tracks files in the system. It manages a database of metadata for each data file in the Enstore system. The metadata includes the file's name, its unique identifier (the bit file ID, or bfid, that the FC itself assigns to each new file), the volume on which it resides, and so on. The FC provides the command:

```
% enstore file --[<option>]
```

which comes with a host of options that allow the user to request specific information about data files. .

5.2 Volume Clerk

The Volume Clerk (VC) is a server that stores and administers storage volume (tape) information. The VC provides the command:

```
% enstore vol --[<option>]
```

which comes with a host of options that allow the user to request specific information about individual volumes. .

5.3 Library and Virtual Library

A Library in Enstore is comprised of both the physical media and a robot arm used to mount the media in attached drives. An Enstore library is typically called a robot. A library/robot interfaces to software that controls the robot arm (the Media Changer, see section 5.6). Each library can contain a variety of media types and employ different types of media drives.

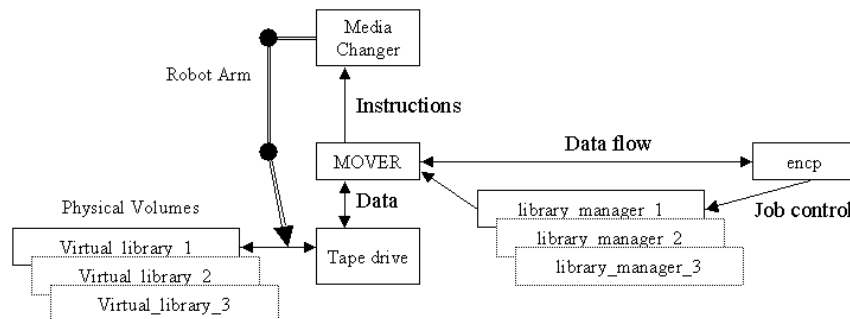
A Virtual Library (VL) is a subset of an Enstore library. It can contain one and only one type of media. It has a Library Manager (LM) associated with it that controls it, as described in section 5.4 *Library Manager*. There may be more than one LM-VL pair for each media type, but not vice-versa. For example, Enstore would divide an STK Powderhorn library holding 50, 20 and 10 GB media into at least three VLs.

5.4 Library Manager

A Library Manager (LM) is a server which is bound to a single Virtual Library (VL), and controls what happens within that VL. There may be many LM-VL pairs in an Enstore system. An LM receives requests for file read and write jobs from the user via **encp**, stores these unassigned requests in a queue, prioritizes them, and dispenses the jobs to a Mover for actual data transfer to and from its VL.

5.5 Mover

A Mover (MV) is a server responsible for efficient data transfer between the **encp** process and a single, assigned media drive. The Mover receives jobs from a Library Manager (LM), and gives instructions to the Media Changer (MC) (described in section 5.6) servicing the Mover's assigned drive.



A Mover is bound to one drive, and a drive may service only one Virtual Library at a time. Each Mover correspondingly services only the LM associated with its drive's VL. However, the actual LM-VL pair serviced at a given moment by the Mover-drive combination may change. The Mover's configuration can list multiple LMs; it requires a file edit by the administrator of your Enstore system to change from one LM-VL pair to another. Allowing flexible LM assignment has two benefits:

- First, since an LM handles only one type of media, a drive which can handle multiple types of media (i.e., different capacity media) can be shared without a static partitioning of the system.
- Secondly, suppose user groups A and B want to share the capacity of a VL, in which half the tapes belong to group A and the other half to group B. You want to guarantee that groups A and B each get one third of the tape drives, and that the last third is shared. To do this, your administrator can configure the Movers to partition resources in the VL, and assign an LM to each type of use.

5.6 Media Changer

The Media Changer (MC) mounts and dismounts the media into and out of drives according to requests from the Movers. One MC can serve multiple drives and thus multiple VLs (the image in section 5.5 *Mover* shows an MC associated with only one drive). When the drives are in the robot, the MC is the interface to the robotic software.

5.7 Configuration Server

The Configuration Server (CS) maintains and distributes the information about Enstore system configuration, such as the location and parameters of each Enstore component and/or server. At startup, each server asks the CS for the information it needs (e.g., the location of any other server with which it must communicate). New configurations can be loaded into the CS without disturbing the current running system.

5.8 Inquisitor

The Inquisitor monitors the Enstore servers, obtains information from them, and creates reports at regular intervals that can be viewed on the web under <http://hppc.fnal.gov/enstore/>. See section 5.11 *Event Relay* for an illustration of an Inquisitor task. The reports created by the Inquisitor include **Enstore Server Status** (section 6.4), **Encp History** (section 6.10), **Enstore Configuration** (section 6.11), and **Enstore Log Files** (section 6.13).

If the Inquisitor goes down (indicated by a red ball next to *Inquisitor* on the **System-At-A-Glance** web page described in section 6.3), data can still be transferred via **encp**. However, the information on the **Enstore Server Status** page (described in section 6.4) doesn't update in this case.

5.9 Alarm Server

The Alarm Server (AS) maintains a record of alarms raised by other servers, and creates a report available online and described in section 6.12 *Enstore Active Alarms*. Since Enstore attempts error recovery whenever possible, it is expected that raised alarms will need human intervention to correct the problem. The AS compares each newly raised alarm with the previously raised ones in order to prevent raising the same alarm more than once.

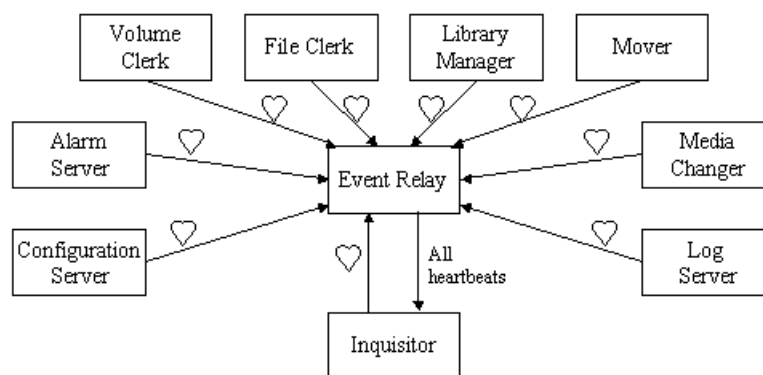
5.10 Log Server

The Log Server (LS) receives messages from other processes and logs them into formatted log files available online and described in section 6.13 *Enstore Log Files*. These messages are transactional records. Log files are labeled by date. Every night at midnight, the currently opened log file gets closed and another one is opened.

5.11 Event Relay

The Event Relay (ER) is a server that acts as a repository of messages. All the Enstore servers send messages to the ER. Some of these servers “subscribe” to the ER in order to have messages of particular types forwarded to them.

For example, the ER periodically receives “I’m alive” messages (called *heartbeats*) from the other servers in the system. The Inquisitor (section 5.8) subscribes to the heartbeat messages, so the ER forwards these messages to it. This is illustrated in the image below:



If the ER goes down (indicated by a red ball next to *Event Relay* on the **System-At-A-Glance** web page described in section 6.3), data can still be transferred via **encp**. However, the information on **Enstore Server Status** (see section 6.4) and the other web pages described in Chapter 6: *Monitoring Enstore on the Web* may not be accurate in this case.

